

Attorney Docket No. 081862.P160
Express Mail No.: EL236840515US

UNITED STATES PATENT APPLICATION

FOR

RE-ROUTING CONNECTIONS USING REDUNDANT PATH
CONNECTIONS AND LOOPBACKS

INVENTORS:

John Ellis
Jan Medved

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 Wilshire Blvd., 7th Floor
Los Angeles, CA 90025-1026
(714) 557-3800

RE-ROUTING CONNECTIONS USING REDUNDANT PATH CONNECTIONS AND LOOPBACKS

Field of the Invention

This invention relates to computer networks. In particular, the invention
5 relates to virtual path channel protection switching.

The Background of the Invention

One way of scaling large networks is to aggregate the channel connections carrying user traffic into the path channels. A failure of such aggregating path channels results in a need to re-route a large number of connections. Such a failure, referred to as channel connection outage, may lead to undesirable performance and reduces customers Quality of Service (QoS).

Traditional techniques to address the connection outage include the use of multiple channel connections. However, the problem with using multiple channel connections is that significant time is required to re-route the user 15 channel connections to a new path. Typically, the channel connections have to be removed and then re-established on the new path. When the number of connections is high, as is typical in a network, the time spent to re-establish the new path may reach several minutes. For service-oriented applications such as high availability applications, this outage time is undesirable.

SUMMARY OF THE INVENTION

A method and apparatus are described for re-routing user connections between first and second nodes in a network switch. A loop-back path provides connectivity between the first and second nodes. The first node has a primary connection and a secondary connection. The primary connection carries the user connections during a normal mode. A switching element is coupled to the loop-back path and the first node to switch the connectivity from the primary connection to the secondary connection when there is a failure condition at the primary connection.

5 Other features and advantages of the invention will be apparent from the detailed description and drawings provided herein.

10

02-142352-1 02-142352-2

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings in which like references indicated similar elements which:

5 Figure 1 shows a system in which one embodiment of the invention can be practiced.

Figure 2 shows a virtual path loop back circuit according to one embodiment of the invention.

10 Figure 3 shows a computer system to perform switching for the loopback path according to one embodiment of the invention.

Figure 4 shows a flowchart for a process to re-route connections.

SEARCHED INDEXED SERIALIZED FILED

DETAILED DESCRIPTION

A method and apparatus are described for re-routing user connections between first and second nodes in a network switch to reduce connection outage in the presence of a failure. A loop-back path provides connectivity between the 5 first and second nodes. The first node has a primary connection and a secondary connection. The primary connection carries the user connections during a normal mode. A switching element is coupled to the loop-back path and the first node to switch the connectivity from the primary connection to the secondary connection when there is a failure at the primary connection. The 10 loop-back path can be either a physical connection or a logical connection. The failure of the primary connection is detected by a network monitor, such as the operations, administration and maintenance (OAM) monitor or a call release procedure. In one embodiment, a connectivity monitor keeps track of a connectivity status between the first and second nodes. The connectivity status 15 indicates end-to-end connections between the first and second nodes. The switching element switches the connectivity based on the connectivity status provided by the connectivity monitor. The secondary connection does not carry user connections during the normal mode, and therefore, does not use the network bandwidth. Typically, the primary and secondary connections have 20 equal connection capacity.

In an embodiment, the network switch is an asynchronous transfer mode (ATM) switch where the primary and secondary connections correspond to a virtual path connection (VPC) in the ATM switch. In such an ATM network, the network monitor may be an operations, administration, and maintenance (OAM) 25 monitor or a call release procedure.

The advantages of the present invention include reducing outages on user SPVC's and SVC's and simplifying topology of the network access.

In the following description, the notation # after a signal name indicates that the signal has an active LOW level, i.e., the signal is asserted when it has a logic LOW level. The description refers to the ATM model and the PCI bus as an interface example. It is contemplated that the technique is applicable to other models, buses, or network architectures with similar characteristics.

Figure 1 shows a system 100 in which one embodiment of the invention can be practiced. The system 100 includes nodes 110, 120, 132, 134, 136, 138, 140, and 150. The system 100 represents a switching network in a communication network. In one embodiment, the network is Asynchronous Transfer Mode (ATM) network.

Each of the nodes 110, 120, 132, 134, 136, 138, 140, and 150 is a switch that performs switching functions to provide connectivity. In one embodiment, each of the nodes is an ATM switch. The nodes 132, 134, 136, and 138 form redundant switched virtual paths (SVP's) over a private network to network interface (PNNI) core. The nodes 110 and 120, 140 and 150 form the edge connections. Over the network, there is a network monitor, such as the operations, administration and maintenance (OAM) monitor or a call release procedure to monitor the virtual path (VP) state.

The nodes 120 and 140 have virtual path (VP) loop back circuits that allow switching of user connections from a primary connection to a secondary connection in the PNNI core. For example, the nodes 132 and 134 may form a primary connection and the nodes 138 and 136 may form a secondary connection. The VP loop back circuits (nodes 120 and 140) receive information on the VP state to determine if there is a failure or a switching condition on the primary connection. If there is a failure, or there is some traffic condition that requires switching, the VP loop back circuits 120 and 140 switch the user connections from the primary connection to the secondary connection. During

the normal mode, the secondary connection does not use the network bandwidth. Therefore the bandwidth is not used for the secondary connection. Typically the primary and the secondary connections have equal capacity so that the user connections are not lost during the switch-over.

5 Figure 2 shows a virtual path loop back circuit 120 according to one embodiment of the invention. The VP loop back circuit 120 includes a loop back path 210, a switching element 220, ports 222 and 224, and a re-route handler 230. The VP loop back circuit 120 provides a loop-back path to provide connectivity between the first node 140 and the second node 110 using a connectivity monitor 10 240. Note that the virtual path loop back 140 also has a similar structure to provide connectivity between the node 120 and the node 150.

15 There are two connections between node 120 (the VP loop back circuit) and node 140: a primary connection and a secondary connection. The primary connection carries the user connections during a normal mode while the secondary connection carries the user connections during a failure mode. The 20 secondary connection is set up or configured at the initialization phase to have essentially the same connection characteristics as the primary connection.

The loop back path 210 includes connection paths, either physical or logical, to connect the node 140 to the node 110. The switching element 220 is coupled to the loop-back path 210 and the node 140 via ports 222 and 224 to switch the connectivity from the primary connection to the secondary connection when there is a failure at the primary connection. The port 222 is connected to the primary connection and the port 224 is connected to the secondary connection. In one embodiment, the switching element 220 is an ATM switching 25 fabric and the ports 222 and 224 are interface ports on an ATM line card.

The re-route handler 230 receives the connection status on the primary connection and controls the switching element 220 to switch from the primary

node to the secondary node in case of failure. The re-route handler 230 receives the connectivity status from the connectivity monitor 240. The connectivity monitor 240 may be part of an OAM procedure or a call release procedure. The connectivity monitor 240 keeps track of a connectivity status between nodes 140 and node 110. The connectivity status indicates end-to-end connections between the nodes 140 and node 110.

Figure 3 shows a computer system 300 to perform switching for the loop-back path according to one embodiment of the invention. The computer system 300 may be used as part of an ATM switch, a host machine, a workstation, a LAN, and any other system or subsystem connected to the network. The computer system 300 include a processor 305, a host bus 310, a host bridge chipset 320, a system memory 330, a peripheral bus 340, a mass storage device 350, a network interface device 355, and K peripheral devices 360₁ to 360_K.

The processor 305 represents a central processing unit of any type of architecture, such as an embedded micro-controller, a general-purpose processor, a digital signal processor, complex instruction set computers (CISC), reduced instruction set computers (RISC), very long instruction word (VLIW), explicitly parallel instruction set computing (EPIC), or hybrid architecture. The invention could be implemented in a multi-processor or single processor computer system.

The host bridge chipset 320 includes a number of interface circuits to allow the host processor 305 access to the system memory 330 and the peripheral bus 340. The host bridge chipset 320 may include a memory controller, a bus interface circuit, and an I/O controller. The memory controller provides an interface to the system memory 330. The I/O controller provides control of I/O functions.

The system memory 330 represents one or more mechanisms for storing information. For example, the system memory 330 may include non-volatile or volatile memories. Examples of these memories include flash memory, read only memory (ROM), or random access memory (RAM). The system memory 5 330 contains a program 332, a data storage 334, and the re-route handler 230 as shown in Figure 2. The re-route handler 230 is a program code or functions to control the switching element 220 in Figure 2 to switch the user connections from a primary connection to a secondary connection based on a connection status such as a failure on the primary connection. Of course, the system memory 330 10 preferably contains additional software (not shown), which is not necessary to understanding the invention.

15 The peripheral bus ³⁴⁰ ~~360~~ provides bus interface to the mass storage device 350, the network interface 355, and the peripheral devices 360₁ to 360_K. In one embodiment, the peripheral bus 360 is the peripheral component interconnect (PCI) bus.

The mass storage device 350 include CD ROM, floppy diskettes, and hard drives. The mass storage device 350 stores non-volatile information such as programs or data. The mass storage device 350 provides a mechanism to read machine-readable media. When implemented in software, the elements of the 20 present invention are essentially the code segments to perform the necessary tasks. The program or code segments can be stored in a processor readable medium or transmitted by a computer data signal embodied in a carrier wave, or a signal modulated by a carrier, over a transmission medium. The "processor readable medium" may include any medium that can store or transfer 25 information. Examples of the processor readable medium include an electronic circuit, a semiconductor memory device, a ROM, a flash memory, an erasable ROM (EROM), a floppy diskette, a compact disk CD-ROM, an optical disk, a

hard disk, a fiber optic medium, a radio frequency (RF) link, etc. The computer data signal may include any signal that can propagate over a transmission medium such as electronic network channels, optical fibers, air, electromagnetic, RF links, etc. The code segments may be downloaded via computer networks such as the Internet, Intranet, etc.

5 such as the Internet, Intranet, etc.

The network interface device 355 provides interface to a network such as ATM, LAN, WAN, etc. In one embodiment, the network interface device 355 provides interface to an ATM switching fabric acting as the switching element 220 shown in Figure 2. The peripheral devices 360₁ to 360_k may include an audio device, a multimedia device, a modem, a printer controller, etc.

Figure 4.1. Summary of the 1994-1995 survey of the U.S. population aged 12 and older.

Upon START, the process 400 connects the first node to the second node via a virtual loop back path on a primary connection during a normal mode (Block 410). Then, the process 400 sets up or configures a secondary connection between the first node and the second node (Block 420). Then the process 400 receives a connectivity status from a network monitor (Block 430).

Next, the process 400 determines if the connectivity status indicates that there is a failure condition at the primary connection (Block 440). If there is no failure condition, the process 400 is terminated. If there is a failure condition, 20 the process 400 switches the connectivity from the primary connection to the secondary connection (Block 450). Then the process 400 is terminated.

A technique has been described to re-route user connections between first and second nodes in a network switch. A loop-back path provides connectivity between the first and second nodes. The first node has a primary connection and a secondary connection. The primary connection carries the user connections during a normal mode. A switching element is coupled to the loop-back path and the first node to switch the connectivity from the primary

connection to the secondary connection when there is a failure condition at the primary connection.

In the foregoing specification, the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

SEARCHED - INDEXED - SERIALIZED - FILED